

Tabu search—Uncharted domains

Fred Glover

© Springer Science + Business Media, LLC 2006

The setting could have come from a Hollywood science fiction movie. A dozen figures, a handful in uniforms of the U.S. Strategic Air Command (SAC) and a slightly larger contingent variously in business suits and in shirtsleeves, were gathered around a large mainframe computer whose blinking lights signaled a run in progress. The computer run — on an early-70s behemoth that represented the state-of-the-art of the period—was seeking a response to a scenario of a first wave nuclear strike on the U.S..

Each of us present had a security clearance at the level of Secret or higher. Mine, a lowly Secret clearance, had come via an engagement with the Defense Communications Agency at the other-worldly SAC installation buried beneath Cheyenne Mountain. My involvement in the present project was the result of a labyrinthine trail that had wended its way through my earlier associations with research groups at Carnegie Mellon University and at the University of California, Berkeley to reach me at my present post at the University of Colorado.

The computer scenario currently in progress was engaged in testing a new type of search procedure applied to the objective of discovering an appropriate counter to the hypothetical nuclear strike. The conditions surrounding this objective, which hinged on satisfactorily meeting a collection of prioritized sub-goals, were embodied in a complex combinatorial optimization model. For the search routine to be successful, an effective retaliatory response had to be identified in less than 5 minutes.

Experience suggested such an undertaking might not be easy. Previous efforts, utilizing a diverse array of approaches implemented by the “shirtsleeves crowd” surrounding me — a computer science and OR group from a California Think Tank organization engaged by SAC for this project — had been unable to generate an effective response in under 47 minutes, and the resulting solution had been deemed marginal. A concerted initiative to remedy this outcome had, reasonably or unreasonably, wound up on my doorstep and confronted me with the challenge of designing an approach that might do better. The California Think Tank group did not conceal their doubts about the likelihood of my pulling this off, and had openly

F. Glover (✉)
Leeds School of Business, University of Colorado, Boulder, CO 80304
e-mail: fred.glover@colorado.edu

conveyed their skepticism to the SAC officers present. The atmosphere was decidedly less than encouraging.

At the signal that the run had been initiated, I was unable to refrain from watching the clock on the wall above the computer. Unexpectedly, far too soon — before the clock had advanced by more than a third of a minute — the lights on the computer display froze. Evidently something had gone wrong. A few snorts on either side of me communicated the message “suspicions confirmed.” Some moments later, however, the chattering of a printer gave evidence that a verification routine had given a clean bill of health to the program execution, and the results were soon common knowledge. The run had been a success, finding a solution significantly better than the best previously seen — one that more than met the criteria of effectiveness on all counts — while consuming at least two orders of magnitude less CPU time.

This occasion did not in reality mark the first time the search procedure implemented for the SAC challenge—the one we now call *tabu search*—had been put to use. Two years earlier a University of Colorado colleague, Claude McMillan, and I had launched a company to provide solution software for workforce planning applications, and the underlying strategies had initially been put through their paces in that setting.

But the spark that ignited the ideas at the heart of the approach came from an incident some years earlier still. In a project for a graduate artificial intelligence course in the early 1960s, my classmates and I were challenged to devise a computer model to emulate some aspect of problem solving. I had been fascinated by the area of integer programming (discrete variable optimization), and was curious to discover how my friends might go about solving a problem in this domain—under conditions where they were not told the nature of formal solution approaches, but were allowed to proceed solely on the basis of their own ingenuity. To pursue the investigation, I devised an experiment that allowed them to work on numerical puzzles that were integer programs in disguise.

As it happened, my friends exhibited intriguing similarities in the way they undertook to find the best solutions to the puzzles. They all employed search patterns that systematically avoided reintroducing elements of constructions recently tried, except where this produced an outcome that appeared too good to pass up. A few of my friends played this game very well, and improved their strategy by isolating attractive possibilities along the way, which they revisited when a particular vein of exploration became unpromising. After intervals where no gains materialized, these efforts were sometimes amplified to bring about more radical changes, giving greater influence to criteria that grew out of obstacles previously encountered.

I was impressed by the effectiveness of the strategies used by the better problem-solvers, and I tried to embody an approximation of their approaches in a computer program for the assigned AI project. In spite of my limitations in translating the relevant processes into computer code, the resulting program performed tolerably well, and in some cases surprisingly well — though it did not measure up to the performance of a couple of my cleverer friends. I made a promise to myself to come back and explore these ideas again someday. (A fuller description of this experiment can be found in Glover (1998)).

Some ten years intervened, involving a series of captivating side excursions into various realms of combinatorial optimization, graph theory and network optimization, before I finally made good on my promise. Incorporating adaptations to fit the context of applications currently encountered gave rise to the computer implementations for the workforce planning and SAC problems. The successes of these applications were striking in view of the complexity of the problems addressed, and persuaded me the underlying ideas had broader application.

Perhaps contrary to what may be imagined, my first efforts to publish these ideas were a bit disheartening. The prevailing opinion in those days regarded heuristic search procedures to lie outside the proper focus of journals devoted to optimization. Lacking a formal convergence guarantee, they were considered guilty of an unpardonable sin of omission. I resorted to pleading extenuating circumstances, and finally, after a prolonged (though cordial) skirmish with referees, I was able to secure acceptance for a substantially condensed version of the material that found its way into print in the late 70s.¹

The immediate aftermath was to provide a vindication of the referees' views. The response to the article was virtually non-existent. Once more I turned my attention for a period to other pursuits, and then revisited the basic ideas a few years later to produce methods for scheduling and space planning applications. Each of these reinforced the outcomes of earlier experiences and met with gratifying success. The resulting publications, on the other hand, met with the same fate as their predecessor. They generated no response.

It was not until a survey paper published in 1986, which referred to the approach by the name tabu search,² that the reception of the underlying ideas began to change. Looking back, however, it is evident that the real reason for the changed reception had little or nothing to do with the new nomenclature, and everything to do with a handful of researchers whose work brought the method to the attention of a broader community. My role was to light a match, and theirs was to build a fire.³

Today the research devoted to tabu search and its variations is a virtual maelstrom of activity, going far beyond anything I could have imagined those many years ago when undertaking the first fledgling implementations and struggling to expose the core ideas. But if the road leading to the current profusion of tabu search applications has been somewhat rocky,⁴ it must be admitted that the reasons are not entirely due to accidents of timing or to early resistance to methods lacking a guarantee of convergence. The initial reception of TS was undoubtedly influenced by the fact that its underlying philosophy clashed with the orientation of other search approaches that dominated the scene at the time.

Departures from familiar terrain

As a counter-current of opinion was soon to attest, several elements of tabu search run against the grain of popular views about the best way to explore a complex space. Researchers from

¹ The paper (Glover, 1977) introduced the basic strategic oscillation approach and associated "strong determination" and "consistency" notions involving frequency memory and statistical analysis. (Recency memory, sometimes considered a primary feature of TS, was illustrated as a way to implement strategic oscillation).

² I have sometimes had misgivings about this name, although perhaps the choice was salutary for subliminal reasons, at least for those with a sense of humor. As friends have pointed out, the label seems to suggest a quest for forbidden fruit, or perhaps a private parlor game for consenting adults.

³ Early instrumental work was led by Dominique de Werra, Alain Hertz, Eric Taillard and Frederic Semet at the Swiss Federal Institute of Technology. A significant role was played by Pierre Hansen, then at Rutgers, who independently formulated basic TS notions and together with Brigitte Jaumard likewise contributed to early implementations. Shortly afterward, a University of Montreal team led by Michel Gendreau, Gilbert Laporte, Jean-Yves Potvin, Teodor Crainic and Jacques Ferland provided innovations that had a significant additional impact. Impossible to overlook, of course, are the contributions of Manuel Laguna, my close collaborator at the University of Colorado. The roll call of researchers who are now extending the frontiers of the method includes many eminent members of the optimization community, and is unfortunately too long to detail here.

⁴ As a partial indication of the level of current activity, a Google web search on the phrase "tabu search" returns about 500,000 pages. A useful source of information and applications can be found on the website www.tabusearch.net, created by Cesar Rego. Vignettes involving seminal applications can also be found in the book *Tabu Search* (Glover and Laguna, 1997).

other backgrounds have shown no hesitation in suggesting that the successes of TS are counter to reasonable expectation, violating principles everyone knows to be the key to finding good solutions to hard problems.⁵

This disposition to reject certain aspects of tabu search sometimes strikes me as bewildering, because it seems to imply a rejection of strategies that people use in their own efforts to solve problems. The folklore of search has inclined us to find special virtue in “memoryless methods” (such as simulated annealing) or “rigid memory methods” (such as branch and bound), or even in methods that hover in a twilight zone of loosely transmitted “inheritance memory” (such as various forms of evolutionary approaches). Yet, in curious contrast, an eminent body of traditional wisdom inclines us to see little of interest in designs anchored to adaptive memory and associated strategies for exploiting it.

The TS focus on adaptive memory admittedly poses challenges that many researchers may prefer to sidestep. But these challenges may also help to define the frontiers of research into problem-solving processes that seek to incorporate abilities analogous to our own. Perhaps the human impulse to avoid acknowledging that our “modern methods” are still somewhat primitive steers us away from ideas that expose the yawning chasms of our ignorance. In any case, the blueprints drawn up for methods that do not venture far into the realms of memory and learning seem to capture the devotion of many search practitioners.

Yet the distinctions I have just alluded to do not incite the degree of resistance to tabu search they once did. Views regarding strategies that are worthy of consideration have changed considerably during the past dozen years. (See, for example, Blum and Roli, 2003; Crainic and Toulouse, 2003; Glover, Laguna, and Marti, 2000). Perhaps it would unduly credit the influence of TS to suggest that it has had a major role in this transformation, but there is no hiding the fact that a variety of metaheuristic approaches have begun adopting TS-based strategies that were once considered outside the dominion of accepted principles. A few articles have even surfaced to endorse the message that adaptive memory is the hallmark of any method that claims to be intelligent, marking a significant departure from the conceptions of a few years ago.⁶

Randomization: For better or worse

Apart from the issue of adaptive memory, a topic that conspicuously separates TS from other approaches concerns the application of random choice. It is impossible to look very far in the metaheuristic literature without becoming aware that it is engaged in a love affair with randomization. This perhaps stems in part from the influence of quantum physics, which endows randomization with a mystique derived from its presence in the behavior of elementary

⁵ Although it would seem better to have a counter-intuitive method that works well than to have a thoroughly “sensible” method that works badly, there is a price to pay for stepping off the beaten path. A number of papers have appeared that have replaced basic TS ideas and strategies with others the authors find more natural. The resulting transmuted methods may not perform in a dazzling fashion, but unfortunately they often perform well enough to get published under the tabu search label—tending to lead other researchers to misinterpret the method and to overlook ideas that offer greater promise.

⁶ Sometimes this message is still improperly conveyed, however, by intimating that adaptive memory should preclude strongly regimented memory as used in tree search. By the tabu search orientation, different kinds of memory are appropriate for different kinds of tasks. For example, the TS *adaptive memory projection* methods are organized to generate sub-problems to be handled by processes, and associated memory structures, that involve varying degrees of regimentation. At the strongly regimented end of the scale, special instances of TS recency and frequency memory have been shown to yield novel variants of tree search, while others yield “non-duplicating” and “rarely duplicating” search patterns that are not readily classifiable.

particles, as expressed in the Heisenberg Uncertainty Principle – fostering the impression that randomization must somehow be a key to the function of the universe. Einstein’s belief that God does not roll dice is out of favor, and many find a special enchantment in miraculous events where blind purposelessness creates useful order. (We are less often disposed to notice that this way of producing order requires an extravagant use of time, and that order, once created, is considerably more effective than randomization in creating still higher order).

Our “scientific” reports of experiments with nature reflect our fascination with the role of chance. When apparently chaotic fluctuations are brought under control by random perturbations, we seize upon the random element as the key, while downplaying the importance of attendant restrictions on the setting in which randomization operates. The diligently concealed message is that under appropriate controls, *perturbation* is effective for creating outcomes of desired forms. If the system and attendant controls are sufficiently constrained, perturbation works even when random, but a significant leap is required to conclude that randomization is preferred to intelligent design. (Instead of accentuating differences between workable and unworkable kinds of perturbation, in our quest to mold the universe to match our mystique we often portray the source of useful outcomes to be randomization in itself).⁷

The tabu search orientation evidently contrasts with this perspective. Many of the fundamental TS strategies do not require randomization for their execution. While TS principles do not exclude recourse to randomization, neither do they embrace it as essential. In general, variants of TS that incorporate randomized elements operate in a context where variation is tightly controlled by strategy. From this point of view, if one is going to play with dice, it’s highly preferable to use dice that are *loaded*. (Perhaps even Einstein’s God might find this a not entirely unreasonable precept).

There is of course a setting where randomization makes good sense. In a game against a shrewd opponent, it can be desirable to make sure one’s behavior exhibits no demonstrable pattern, in order to avoid the chance that the pattern will be detected and turned to the adversary’s advantage. In this case, recourse to randomization is a prudent policy. On the other hand, presuming that our search spaces are not possessed of a devious intelligence bent on thwarting our moves, random behavior fulfills no purpose comparable to its role in game playing. Instead, within the context of search, randomization seems more likely to be a means of outwitting the “player” who uses it. A policy of embarking on a series of random steps would seem one of the best ways of confounding the goal of detecting and exploiting structure.

Granted, systematic search runs the risk of systematic blundering, something none of us is immune to. But with appropriate monitoring, such a form of search also affords an opportunity to isolate and rectify our blunders, and therefore to carry out more effective explorations. (Of course, if the search terrain itself is inherently random, no amount of systematization will do any good. The presence or absence of patterned design makes no difference when all is reduced to blind luck).⁸

⁷ The theme that randomization is an essential underpinning of modern search methods is echoed throughout many segments of the metaheuristic literature. Perhaps this orientation relates to our natural disposition to take comfort in the notion that all can be left up to chance and everything will turn out for the best in the end. Or perhaps randomization has a seductive charm akin to the appeal of a certain kind of romantic encounter — which seems somehow more captivating when marked by an element of capriciousness.

⁸ Even the statement that structured search has little value for a problem whose character is “essentially random” deserves qualification. A sorting problem may involve entirely random data, but still be handled advantageously by a method that is highly systematic. (The “P vs. NP” distinction is relevant, but does not remove the need for additional qualification).

I am tempted to pose a research challenge, as a way of testing the relevance of randomization in search. If a policy of random choice produces a useful sequence of responses, then we may presume that a different randomly generated sequence will likewise prove useful. On the other hand, not all sequences are apt to be equally good if the total number of moves is limited — a relevant consideration if it is important to reach good outcomes before the knell of doomsday, or at least before next summer's vacation. When time matters, we may anticipate that not all randomly generated sequences will provide the same quality of performance.

This observation prompts the following experiment. To determine the relative utility of randomization, we might examine the outcomes of executing some number, say k , of randomly guided searches (each of limited duration) with the goal of finding a value of k such that at least one of the underlying choice sequences will lead to a solution of specified quality. We seek a value for our parameter that is not too large and that is applicable to a large portion of problems from a given class.

Now the evident question comes to mind. If we replace the k randomly generated choice sequences with k systematically generated sequences, can we obtain comparable or better results? Or can the replacement allow us to produce such results for a smaller value of k ? In accordance with the theme of TS diversification strategies, the systematically generated sequences may be equipped to incorporate feedback, allowing a new sequence may be influenced by the outcomes of preceding sequences. (There is an issue of time invested in the systematic approach, so that the value of k must be adjusted to account for this).

The outcome of such an experiment of course depends on our skills in designing systematic choice sequences, and also depends on the types of problems we confront. (My biases lead me to suspect these problems are often susceptible to the ploys of our ingenuity, in part because we are the ones who define these problems and hence to some degree determine their form).

It must be acknowledged that the experiment I am advocating has been implicit in a variety of TS implementations of the past. But there may be advantages to bringing the relevant factors into more explicit focus. Perhaps this will help us to better understand what is essential and what is superfluous when we speak of "intelligent strategy," and to identify better instances of such strategy over time.

An evident subtext of the preceding commentary is that the discovery of useful non-random patterns is predicated on learning, and learning in turn requires memory. Yet memory poses subtleties that may not be immediately visible on the surface. This motivates us to take a closer look at memory and its implications.

Is memory really a good idea?⁹

The TS orientation that regards memory to be an integral component of intelligent search would not seem unduly radical. From this standpoint, it may at first seem puzzling that some of the methods widely heralded as innovations in artificial intelligence—as applied to optimization — are largely devoid of memory. The reluctance to incorporate non-trivial types of memory in such approaches is not as unreasonable as might be imagined, however, if it is considered relative to a quest for mathematical vindication. Memory, particularly in its adaptive forms, introduces too many degrees of freedom to be treated conveniently in "theorem and proof" developments. Researchers who prefer to restrict consideration to

⁹ The comments of the next few subsections are adapted from material written several years ago, whose relevance is now amplified by emerging views about the appropriate design of search methods. (See Glover and Laguna (1997) and Gendreau (2003)).

processes (and behavior) that can be characterized by rigorous proofs must focus their efforts in other directions.¹⁰

Yet there are more subtle and valid reasons to be wary about the use of memory. Malleable forms of memory entail certain dangers — potential pitfalls that go hand in hand with the ability to provide valuable strategic opportunities. These dangers are the price to be paid for the evolution of “intelligent” mechanisms, including biological mechanisms embodied in a brain.

From an evolutionary standpoint, the emergence of memory may be viewed as posing a challenge comparable to the emergence of oxygen, whose corrosive properties (as evolutionary biologists are fond of telling us) caused considerable destruction until organisms adapted to take advantage of them. Analogous perils may well have been created by the emergence of memory, though today we only see the outcomes that survived and flourished. Characteristically, the blind alleys of poorly designed physical structure are conspicuously imprinted on the fossil record. But the blind alleys of poorly regulated mental adjustment — which may have affected survival in subtler ways—remain invisible to us.

It is noteworthy, however, that we are memory users whose evolutionary line has survived. Since we tend to endow our problem solving schemes with features that reflect our own disposition, such schemes tend to be protected (at least to a degree) from dangers otherwise presented by adaptive memory. Even so, hastily contrived uses of memory can lead to conspicuously undesirable outcomes.

Accordingly, in order to solve complex problems more effectively, TS approaches seek to uncover the potential gains of adaptive memory without being caught in the traps of ill-considered memory designs. This leads to a quest for “integrating principles,” by which alternative forms of memory are appropriately combined with effective strategies for exploiting them. A novel finding is that such principles are sometimes sufficiently potent to yield effective problem solving behavior with negligible reliance on memory. Over a wide range of problem settings, however, strategic use of memory can make dramatic differences in the ability to solve problems.

Fundamental issues

A starting premise of tabu search is that intelligent inhibition plays a critical role in making effective use of memory. This may be conceived as a reflection of an analogous supposition that appropriate forms of inhibition and restraint correspondingly are essential to survival, although we may not always think of such elements as survival tools. (The usual tenet of our culture is that inhibition represents something that must be overcome, rather than something that can provide important advantages when properly utilized).

The connotation of the “tabu” term in tabu search carries an implication, as it does in other domains, of rules that are contextual and subject to change. This type of variability can range from narrowly confined interaction to highly complex coordination. The potential intricacy of managing such variation understandably may pose an obstacle to rules that are too rigidly constrained by the quest for mathematical precision, but there is a reverse danger of seeking to handle complexity by the expedient of simplistic rules, particularly those that rely heavily on randomization as a substitute for identifying strategic relationships.

¹⁰ This remark is to some extent misleading. Within the last few years finite convergence has been established for TS variants based on recency and frequency memory. Notably, however, the approaches that give rise to such proofs lack the richness of the procedures that have proved most effective in practice.

Currently it is fashionable to base the design of search mechanisms on a level of organization represented by primitive organisms. But we may legitimately wonder whether intelligent behavior can be adequately encompassed within physical or biological processes that are distant precursors of the forms of organization embodied in our own brains. If there is value in having the capabilities we call human, then it seems questionable to aspire to mimic something less.

There is of course no reason to limit consideration to forms of intelligence that match our own. Evolution presumably may have honed our skills to handle problems that have typically presented themselves in our surroundings. Our prowess may be less impressive for problems confronted in other settings, including problems created as a result of our own technology. A leading goal of TS research is to identify memory and strategy combinations that have merit in a wide range of contexts, not restricted to those we have commonly encountered by the accident of history. If this pursuit may yield insights into different types of intelligence, that would be a welcome bonus. Conceivably, by this perspective, the field of memory-based search methods may have something useful to contribute to the field of cognitive behavior. Up to now, the complementary realms of *search* and *psychology* have been largely isolated from each other. This situation may change, however, as findings about the connections between adaptive memory processes and improved problem solving become systematized.

Metaphors of nature

The types of search we find most reasonable to use are often intimately bound to the images we rely upon to portray these approaches. A popular thrust of many research initiatives, and especially of publications designed to catch the public eye, is to associate various methods with processes found in nature. This trend embodies a wave of “New Romanticism,” reminiscent of the Romanticism of the 18th and 19th centuries (distinguished by their preoccupation with Nature with a capital “N”). The current fascination with natural phenomena as a foundation for problem-solving methods is undoubtedly fueled by our sense of mystery concerning the ability of such phenomena to generate outcomes that are still far beyond our comprehension. However, the New Romanticism goes farther, to suggest that by mimicking the rules we imagine to operate in nature (especially “rudimentary” processes of nature) we will similarly be able to produce remarkable outcomes.

Models of nature that are relied upon for such inspiration are ubiquitous, and it is easy to conjure up examples whose metaphorical possibilities have not yet been tapped. To take an excursion in the lighter side of such possibilities (though not too far from the lanes currently traveled), we may observe that a beehive offers a notable example of a system that possesses problem solving abilities. Bees produce hives of exceptional quality and complexity, coordinate diverse tasks among different types of individuals, perform spatial navigation, and communicate via multiple media. (It is perhaps surprising in retrospect that the behavior of bees has not been selected as a basis for one of the “new” problem solving methods).¹¹

There is evidently no need to confine our attention to sentient creatures for analogies that inspire templates for effective problem solving. The root system of a tree, for example, provides an intriguing model for parallel computation. In order to find moisture and nutrients (analogous to a quest for “solutions”), roots distribute themselves across different regions,

¹¹ The preceding comment, taken from Chapter 1 of Glover and Laguna (1997), has an amusing sequel. Websites to promote methods based on so-called “swarm intelligence” have recently appeared that display the image of bees circling a hive.

sending out probes that multiply or atrophy according to the efficacy of their progress. The paths of such a system may cross, as different channels prove promising by virtue of the regions in which they lie and also according to the directions in which they are explored. Obstacles are effectively skirted, or over time are surmounted by longer-range strategies — as by extending finer probes, which ultimately expand until the medium is broached. There exist some root systems, as in groves of aspen, where roots of one entity can merge with those of another, thus enlarging the potential sources of communication and contact available to each. (Such an interlinked community gives rise to the largest known organisms on the planet).

These analogies to systems in nature invite us to ponder a key question. If we were allowed to place our bets on the probable success of a hive of bees or a grove of aspen, as opposed to that of a group of humans, when confronted with a challenging task that requires intelligence and the ability to learn from the past, how would we wager? Undoubtedly we would be drawn to reflect that our goals and problem structures may often be different than those to which “natural processes” apply. In addition, we ourselves—as products of a rather special and extended chain of natural developments—may incorporate capabilities not present in the processes that produced us.

Metaphors of nature have a place. They appear chiefly to be useful for spurring ideas to launch the first phases of an investigation. As long as care is taken to prevent such metaphors from cutting off lines of inquiry beyond their scope, they provide a means for “dressing up” the descriptions of various metaheuristics in a way that appeals to our instinct to draw parallels between simple phenomena and abstract designs.

Invoking such parallels may sometimes appear to embody a primitive mysticism, akin to chanting about campfires in the night, but it gives us a foundation for connecting the new to the old, and for injecting passion into our quests. It is up to prudence to determine when the symbolism of the New Romanticism obscures rather than illuminates the pathway to improved understanding. Within the realm of metaheuristic design, there is a great deal we have yet to learn. The issue of whether the analogies that underlie some of our models may limit or enhance our access to further discovery deserves careful reflection.

Conclusions

Evidently, this brief sketch of the background and evolution of tabu search, and of certain ideas that lie at its center, is only fragmentary. The fundamental message is that the journey from past to present may have enabled us to solve many kinds of difficult optimization problems more effectively, but we still have a long way to go. Not only does a great deal remain to be learned about tabu search, but it is equally true that very little is yet known about how we ourselves use memory in our problem solving. It is worth stressing again that discoveries about effective uses of memory within our search methods may provide clues about strategies that humans are adept at employing—or may advantageously be taught to employ. The potential links between the areas of heuristic search and psychology are an intriguing concomitant to research now underway and have scarcely been examined. Progress in the design of tabu search methods, and the successful applications of TS that have occurred so far, provide encouragement that such issues are profitable to probe more fully.

References

- Blum, C. and A. Roli. (2003). “Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison.” *ACM Computing Surveys*, 35(3), 268–308.

-
- Crainic, T.G. and M. Toulouse. (2003). "Parallel Strategies for Meta-Heuristics." In G. Kochenberger and F. Glover (Eds.), Chapter 17 of *Handbook of Metaheuristics*. Kluwer Academic Publishers.
- Gendreau, M. (2003). "An Introduction to Tabu Search," in: G. Kochenberger and F. Glover, (Eds.), Chapter 2 of *Handbook of Metaheuristics*, Kluwer Academic Publishers.
- Glover, F. (1977). "Heuristics for Integer Programming Using Surrogate Constraints," *Decision Sciences*, 8(1), 156–166.
- Glover, F. (1986). "Future Paths for Integer Programming and Links to Artificial Intelligence." *Computers and Operations Research*, 13(5), 533–549.
- Glover, F. (1998). "Tabu Search—Wellsprings and Challenges." *European Journal of Operational Research*, 106, 221–225.
- Glover, F. and M. Laguna. (1997). *Tabu Search*. Kluwer Academic Publishers.
- Glover, F., M. Laguna, and R. Marti. (2000). "Fundamentals of Scatter Search and Path Relinking." *Control and Cybernetics*, 29(3), 653–684.